

IN THE U.S. PATENT AND TRADEMARK OFFICE  
Patent Application Transmittal Letter

ASSISTANT COMMISSIONER FOR PATENTS  
Washington, D.C. 20231

Sir:  
Transmitted herewith for filing under 37 CFR 1.53(b) is a(n): ☒ Utility ☐ Design  
☒ original patent application,  
☐ continuation-in-part application

JC841 U.S. PTO  
09/06/00

INVENTOR(S): Zheng Zhang and Manohar Karkal Prabhu

TITLE: CHECKPOINT COMPUTER SYSTEM UTILIZING A FIFO BUFFER TO RE-SYNCHRONIZE AND  
RECOVER THE SYSTEM ON THE DETECTION OF AN ERROR

Enclosed are:

☒ The Declaration and Power of Attorney. ☒ signed ☐ unsigned or partially signed  
☒ 2 sheets of drawings (one set) ☐ Associate Power of Attorney  
☐ Form PTO-1449 ☐ Information Disclosure Statement and Form PTO-1449  
☐ Priority document(s) ☐ Other (fee \$ )

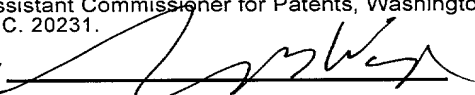
CLAIMS AS FILED BY OTHER THAN A SMALL ENTITY				
(1) FOR	(2) NUMBER FILED	(3) NUMBER EXTRA	(4) RATE	(5) TOTALS
TOTAL CLAIMS	20		X \$18	\$
INDEPENDENT CLAIMS	3		X \$78	\$
ANY MULTIPLE DEPENDENT CLAIMS			\$260	\$
BASIC FEE: Design (31 0 00 ); Utility \$(690.00 )				\$
TOTAL FILING FEE				\$
OTHER FEES				\$
TOTAL CHARGES TO DEPOSIT ACCOUNT				\$ 690

Charge \$ 690 to Deposit Account 08-2025. At any time during the pendency of this application, please charge any fees required or credit any over payment to Deposit Account 08-2025 pursuant to 37 CFR 1.25. Additionally please charge any fees to Deposit Account 08-2025 under 37 CFR 1.16, 1.17, 1.19, 1.20 and 1.21. A duplicate copy of this sheet is enclosed.

"Express Mail" label no. EK 334075425 US

Date of Deposit 9/6/00

I hereby certify that this is being deposited with the United States Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above and is addressed to: Assistant Commissioner for Patents, Washington, D.C. 20231.

By   
Typed Name. Calvin B. Ward

Respectfully submitted,

By 

Calvin Ward  
Attorney/Agent for Applicant(s)  
Reg. No. 30896

Date: 9/6/00

Telephone No.: 925-855-0413

## 5

## 10

## 15

20

25

30

At periodic intervals, the slave computer generates, or receives, a checkpoint signal. The signal may be generated internally via an interrupt timer associated with either of the computers, a hardware or software interrupt generated by either computer, or the signal may be generated by hardware that is external to both CPU's. When the checkpoint signal is received by the slave computer, the slave computer stores the contents of its registers in a predetermined location so that the slave can be returned to the state at the checkpoint. Since no errors have occurred since the last checkpoint, the FIFO buffer is merely dumped, and the next computational cycle is begun.

As noted above, if an error occurs in the master computer system and the master computer system cannot recover from the error, program execution is transferred to the slave computer system, which begins execution from the last checkpoint. The slave computer uses the contents of the FIFO buffer to return the slave memory to the state it had at the end of the last checkpoint. At this point, the contents of the master and slave memories are synchronized as of that checkpoint. The slave computer then loads its registers from the register images stored at the end of the last checkpoint interval and picks up where the master computer left off.

While this system provides a high degree of fault tolerance, it is quite expensive, since it requires a second slave computer.

Broadly, it is the object of the present invention to provide an improved checkpoint computer system.

It is a further object of the present invention to provide a checkpoint computer system that operates on a single computer, thereby eliminating the need for a slave computer.

These and other objects of the present invention will become apparent to those skilled in the art from the following detailed description of the invention and the accompanying drawings.

### **Summary of the Invention**

The present invention is a fault-tolerant computer system having an application memory organized as a plurality of cache lines, each cache line being identified by an address in the memory. A FIFO buffer stores a plurality of such cache lines. The system includes at least one CPU for executing instructions stored in the application memory. The system includes a state memory for storing the contents of the internal registers of the CPU. A checkpoint controller defines a series of repeating checkpoint cycles. The application memory and FIFO buffer are operated under the control of a memory controller. The checkpoint controller also has access to a plurality of registers in the CPU that define the state of that CPU at a point in each checkpoint cycle that is controllable by the checkpoint controller. When the memory controller receives a cache line from the CPU in response to a write command specifying an address A in the application memory at which the cache line is to be stored, a copy of the cache line as stored in the application memory at A is copied into the FIFO buffer upon receiving the first write command specifying A after the start of the current checkpoint cycle. The cache line received in the write command is then used to overwrite the contents of address A in the application memory. At a predetermined point in each checkpoint cycle, the checkpoint controller causes the CPU to cease processing instructions from the application memory; to write back to the memory all dirty cache lines; and to store its internal registers defining the state of the CPU in the state memory. The checkpoint controller empties the contents of the FIFO buffer at the end of each checkpoint cycle if no error has been detected by the end of the checkpoint phase of the cycle. If an error is detected, the contents of the FIFO buffer are read back into the application memory, and the contents of the state memory are read back into the CPU. The system is then restarted after completing any hardware configuration needed for the restart.

### **Brief Description of the Drawings**

Figures 1 and 2 illustrate the various phases of a checkpoint system.

Figure 3 is a block diagram of a multiprocessor computer 10 having a main memory 20 according to the present invention.

### **Detailed Description of the Invention**

The present invention is a checkpoint system that runs on a single computer thereby providing the advantages of the prior art checkpoint systems discussed above without the added cost of maintaining the slave computer system. Referring to Figure 1, there are four phases of operation in a checkpoint system. The first two occur during normal operation and the remaining two occur following the detection of a failure event. The first two phases will be referred to as a checkpoint cycle. The phase of a checkpoint cycle in which normal processing occurs is referred to as the computational phase. During the computational phase, maintenance is done in the background to provide the data needed for system recovery in the event of a fault. The other phase of a checkpoint cycle, referred to as the checkpoint phase, corresponds to the time when the machine is brought into a known state that can be re-created as a starting point in the event of a fault. In this state the cache lines are all clean (i.e. synchronized with respect to main memory). During this phase, the internal state of the processing nodes is also brought to a known state and/or recorded.

The other two phases of operation are expected to occur very rarely. They only occur during the recovery period, after a fault has been detected in the system. The first is the recovery phase. The other is the degraded-operation phase, and follows the recovery phase. Although Figure 2 shows the fault occurring during the computational phase, the fault can occur during the checkpoint phase as well.

The process of checkpointing allows the machine to backtrack to a previously consistent machine state that existed prior to the fault and from which execution can be resumed. This requires that a copy be kept of the state of the machine at the time of completion of the last checkpoint cycle. This is achieved in the present invention by saving the checkpoint state of the machine in a fault-tolerant memory as changes are made to the main memory between checkpoints. In the simplest embodiment of the present invention, each time a line is written back into the main memory, the contents of the main memory that were displaced by that line are recorded in an undo FIFO buffer. If a fault occurs, the state of the memory at the end of the last checkpoint cycle can be reconstructed by reading out the

FIFO and rewriting the main memory from the contents thereof. If no fault occurs during a checkpoint cycle, the contents of the FIFO are merely dumped at the end of the cycle. If the main memory has a number of different nodes, a separate undo FIFO is kept for each memory node to record changes to that memory since the last checkpoint. More sophisticated  
 5 embodiments of the present invention that require smaller FIFO buffers will be discussed in more detail below.

The checkpoint process must store information sufficient to restore the machine to the state that the machine had at the end of the last checkpoint cycle. In addition to storing the  
 10 various registers in the machine, the cache memories must be brought to a reproducible state. The preferred method for accomplishing this is to cleanse the cache memories of dirty lines, which are written back to the main memory, and hence, preserved, if necessary. This write back imposes a load on the bus system during the checkpoint phase. In a distributed  
 15 computing system, all nodes must synchronously perform the checkpoint and complete the checkpoint phase.

The manner in which the present invention operates in a computer having a single main memory may be more easily understood with reference to Figure 3, which is a block  
 20 diagram of a multiprocessor computer 10 having a main memory 20 according to the present invention. Memory 20 stores information for a plurality of processors connected thereto by bus 35. Exemplary processors are shown at 31 and 32. Memory 20 is under the control of controller 41. The timing of the checkpoint cycles is preferably controlled from a checkpoint  
 controller 43 that can be part of controller 41 or a separate unit.

For the purposes of this discussion, it will be assumed that main memory 20 is read  
 25 and written in units of cache lines and each cache line has appended thereto error correcting bits which can be used by controller 41 to detect and correct errors in the cache line. Since error correcting memories of this type are well known in the art they will not be discussed in detail here. In general, these codes rely on the fact that only a small number of the possible  
 30 cache lines are possible if no errors are introduced during the storage and retrieval process. For example, in a typical error-correcting code, each 8-bit data word is transformed into a 16-bit storage word. There are only 256 possible data word values; hence, only 256 of the

possible 65536 storage word values will be used in the absence of errors. When an error occurs, a valid storage word is usually converted to an invalid storage word. The error correcting system then tries to figure out which valid state would have given rise to the detected invalid state if various numbers of bits were altered by the storage operation. The ability of an error correcting code to correct errors is measured by a quantity referred to as the “Hamming Distance” associated with the code. For example, codes with a Hamming Distance of 5 can detect errors resulting from 4 single-bit errors and correct for all possible 2 single-bit errors. A discussion of error correcting codes may be found in ERROR CORRECTING CODES, 2ND EDITION, by Peterson and Weldon, MIT PRESS, 1972, or in PRACTICAL ERROR DESIGN FOR ENGINEERS, by Neil Glover, Data Systems Technology Corp., 1982.

Main memory 20 is divided into an application memory 21 that is used to store the actual cache lines that are delivered to the various processors and an undo-FIFO buffer 22. In the simplest embodiment of the present invention, each time a cache line is written back into application memory 21, the line displaced by the write back is copied into FIFO 22. At the end of each checkpoint cycle, the state of each of the CPUs is written into a predetermined location 25. This data includes the contents of the various registers within the CPUs. In addition, each processor writes back any dirty cache lines to main memory.

Each time a cache line is read back to main memory, it displaces the contents of one of the cache lines in application memory 21. In this simplest embodiment of the present invention, the displaced cache line is written into FIFO 22. If, at the completion of the checkpoint phase, no faults have been detected, the contents of FIFO 22 are dumped by resetting the appropriate pointer.

If a fault is detected, then the contents of FIFO 22 are written back into application memory 21 during the fault recovery phase to return the application memory to the state it had at the end of the last checkpoint cycle. The various CPUs can then be reloaded from CPU status buffer 25 and the program restarted.

While this simple embodiment of the present invention will operate, it requires a large FIFO 22 or a relatively short checkpoint cycle, since FIFO 22 must be capable of storing all of the displaced cache lines during a single checkpoint cycle. Short checkpoint cycles are to be avoided since the checkpoint overhead is more or less fixed, and hence, the efficiency of the computer system decreases with decreasing checkpoint cycle times. In this regard it should be noted that controller 43 can almost always prevent a FIFO buffer overflow by keeping track of the free space in FIFO 22 and starting a new checkpoint cycle if the free space falls below some predetermined limit.

The size of the FIFO buffer can be significantly reduced by keeping track of which cache lines have copies in FIFO 22. Consider the case in which a cache line at address A is written back into the application memory several times during a single checkpoint cycle. The only copy of the cache line in the FIFO buffer that is of any significance is the one displaced by the first write back. Any subsequent lines for this address will be overwritten during recovery leaving this line in the main memory.

In the preferred embodiment of the present invention, a flag is assigned to each cache line in application memory 21. At the beginning of a checkpoint cycle, the flags are all set to a first state indicating that a write back has not occurred for that cache line. When a line is written back, the flag for that line is examined to determine if a write back has already occurred. If a write back has not occurred, the cache line displaced by the write back is copied into FIFO 22, and the flag is set to indicate that a write back has occurred. If a write back has occurred, the displaced cache line is ignored, as the copy needed for recovery of the memory at the address in question is already in FIFO 22.

At the end of a checkpoint cycle, the various flags must all be reset. Since the number of flags is quite large, the resetting process can take a considerable length of time.

Accordingly, in the preferred embodiment of the present invention, two sets of flags are kept as shown at 23 and 24. Each vector has one bit for each cache line in application memory 21.

The vectors alternate between two states. During any given checkpoint cycle, only one of the vectors is utilized to keep track of which cache lines have been written back to the main



memory. The other vector is reset in the background. Hence, at the end of the checkpoint cycle, the unused vector will be reset and the current vector becomes the vector to reset.

As noted above, the cache memories in each of the processors must be synchronized with the main memory at the end of each checkpoint cycle. The corresponding write back activity can saturate bus 35, and hence, lengthen the checkpoint phase of the cycle. In the preferred embodiment of the present invention, controller 41 tracks the lines within the main memory that are to receive a write back during the current checkpoint cycle. For example, if a processor requests an exclusive copy of a cache line from the main memory, controller 41 knows that a write back for that cache line may occur sometime between now and the end of the current checkpoint cycle. If and when the write back occurs, the contents of the memory at that cache line would be read into FIFO 22 if that cache line had not already received a write back. In the preferred embodiment of the present invention, controller 41 does not wait for the write back to actually occur before copying the current contents of application memory 21 at that cache line into FIFO 22 and setting the corresponding flag. If a copy is to be made, i.e., the corresponding cache line flag was not already set, controller 41 copies the contents of the cache line from the main memory into FIFO 22 during the computational phase of the checkpoint cycle, thereby reducing the load during the checkpoint phase.

The above-described embodiments of the present invention utilize a fault-tolerant memory. However, it will be obvious to those skilled in the art from the preceding discussion that embodiments that utilize conventional memory can also be utilized. Such embodiments are not preferred, however, because they cannot recover from memory errors.

While the above-described embodiments have utilized conventional memory for the shared memory, it should be noted that the present invention could be utilized with all forms of memory including disk drives and the like. Embodiments in which the FIFO buffer and/or flag buffers are kept in separate memories can also be practiced without departing from the teachings of the present invention.

The above-described embodiments of the present invention utilize a scheme in which memory is read and written in blocks consisting of cache lines. However, it will be obvious

to those skilled in the art from the preceding discussion that the present invention can be implemented using other types of storage blocks.

As noted above, the present invention stores the state of the machine at the end of each checkpoint cycle. In some hardware configurations such as wave pipelined processors, the status of the machine is not always available in the registers that can be accessed by the checkpoint controller. In such configurations, the checkpoint data on the machine status must be taken at a time when the data is available. Hence, the end of the checkpoint cycle may include the insertion of “no-ops” to flush the pipeline prior to reading the appropriate machine registers so that the machine is in a known state that is ascertainable from the registers.

In the above-described embodiments of the present invention, the machine was merely restarted from the checkpoint data after the detection of a fault. However, it will be obvious to those skilled in the art from the preceding discussion that the machine can also be reconfigured prior to restart. For example, faulty hardware can be eliminated provided the remaining hardware is sufficient to run the program in question.

While the above-described embodiments of the present invention utilized hardware interrupts to initiate a checkpoint cycle, it will be obvious to those skilled in the art from the preceding discussion that software interrupts may also be utilized. Hence, the term “interrupt” as used herein is defined to include both hardware and software interrupts unless the context indicates otherwise.

Various modifications to the present invention will become apparent to those skilled in the art from the foregoing description and accompanying drawings. Accordingly, the present invention is to be limited solely by the scope of the following claims.

## WHAT IS CLAIMED IS:

## 1. A computer system comprising:

5 an application memory organized as a plurality of cache lines, each cache line being identified by an address;

a FIFO buffer for storing a plurality of cache lines;

10 a central processing unit (CPU) for executing instructions stored in said application memory;

a state memory for storing the contents of the internal registers of said CPU;

15 a checkpoint controller for defining a series of repeating checkpoint cycles, said checkpoint controller having access to a plurality of registers in said CPU that define the state of that CPU at a point in each of said checkpoint cycles that is controllable by said checkpoint controller; and

20 a memory controller for operating said application memory and said FIFO buffer, said memory controller receiving a cache line from said CPU in response to a write command specifying an address A in said application memory at which said cache line is to be stored, wherein a copy of said cache line as stored in said application memory at address A, is copied into said FIFO buffer upon receiving the first write command specifying A after the start of  
25 the current checkpoint cycle, said cache line received in said write command replacing the contents of A in said application memory.

2. The computer system of Claim 1 wherein each checkpoint cycle comprises a computational phase and a checkpoint phase and wherein said checkpoint controller during  
30 said checkpoint phase causes said CPU to write back to said memory all dirty cache lines; and to store internal registers defining the state of said CPU in said state memory.

3. The computer system of Claim 2 wherein said checkpoint controller empties the contents of said FIFO buffer at the end of each checkpoint phase if no error has been detected by the end of said checkpoint phase.

5           4. The computer system of Claim 1 wherein said checkpoint controller, in response to a determination that a processing error has occurred, copies the contents of said FIFO buffer into said application memory, causes said CPU to copy the contents of said state memory into said CPU's internal registers, and restarts said computer system.

10           5. The computer system of Claim 4 wherein said checkpoint controller causes said computer system to be reconfigured before restarting said computer system.

15           6. The computer system of Claim 1 wherein said application memory comprises a fault tolerant memory.

20           7. The computer system of Claim 1 further comprises a plurality of reset flags, each reset flag corresponding to one of said cache lines in said application memory and having a first state indicating that the corresponding cache line in said application memory has been re-written since the start of the current checkpoint cycle and a second state indicating that the corresponding cache line has not been re-written since the start of said current checkpoint cycle, the state of said reset flags determining if said memory controller copies the contents of one of said cache lines into said FIFO buffer.

25           8. The computer system of Claim 1 further comprises first and second vectors of flags, each vector being sufficient to provide said plurality of reset flags, said first vector supplying said reset flags in every other checkpoint cycle and said second vector supplying said reset flags in said remaining checkpoint cycles, said vector that is not supplying said reset flags in any given cycle being set to said second state during that cycle.

30           9. The computer system of Claim 1 wherein said memory controller causes a copy of said cache line as stored in said application memory at an address A to be copied into said FIFO buffer upon receiving the first read command specifying A after the start of the current

checkpoint cycle if said read command assigns exclusive ownership of said cache line to said CPU.

	1990	1991	1992	1993	1994	1995	1996	1997	1998	1999	2000	2001	2002	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	2026	2027	2028	2029	2030	2031	2032	2033	2034	2035	2036	2037	2038	2039	2040	2041	2042	2043	2044	2045	2046	2047	2048	2049	2050	2051	2052	2053	2054	2055	2056	2057	2058	2059	2060	2061	2062	2063	2064	2065	2066	2067	2068	2069	2070	2071	2072	2073	2074	2075	2076	2077	2078	2079	2080	2081	2082	2083	2084	2085	2086	2087	2088	2089	2090	2091	2092	2093	2094	2095	2096	2097	2098	2099	2100	2101	2102	2103	2104	2105	2106	2107	2108	2109	2110	2111	2112	2113	2114	2115	2116	2117	2118	2119	2120	2121	2122	2123	2124	2125	2126	2127	2128	2129	2130	2131	2132	2133	2134	2135	2136	2137	2138	2139	2140	2141	2142	2143	2144	2145	2146	2147	2148	2149	2150	2151	2152	2153	2154	2155	2156	2157	2158	2159	2160	2161	2162	2163	2164	2165	2166	2167	2168	2169	2170	2171	2172	2173	2174	2175	2176	2177	2178	2179	2180	2181	2182	2183	2184	2185	2186	2187	2188	2189	2190	2191	2192	2193	2194	2195	2196	2197	2198	2199	2200	2201	2202	2203	2204	2205	2206	2207	2208	2209	2210	2211	2212	2213	2214	2215	2216	2217	2218	2219	2220	2221	2222	2223	2224	2225	2226	2227	2228	2229	2230	2231	2232	2233	2234	2235	2236	2237	2238	2239	2240	2241	2242	2243	2244	2245	2246	2247	2248	2249	2250	2251	2252	2253	2254	2255	2256	2257	2258	2259	2260	2261	2262	2263	2264	2265	2266	2267	2268	2269	2270	2271	2272	2273	2274	2275	2276	2277	2278	2279	2280	2281	2282	2283	2284	2285	2286	2287	2288	2289	2290	2291	2292	2293	2294	2295	2296	2297	2298	2299	2300	2301	2302	2303	2304	2305	2306	2307	2308	2309	2310	2311	2312	2313	2314	2315	2316	2317	2318	2319	2320	2321	2322	2323	2324	2325	2326	2327	2328	2329	2330	2331	2332	2333	2334	2335	2336	2337	2338	2339	2340	2341	2342	2343	2344	2345	2346	2347	2348	2349	2350	2351	2352	2353	2354	2355	2356	2357	2358	2359	2360	2361	2362	2363	2364	2365	2366	2367	2368	2369	2370	2371	2372	2373	2374	2375	2376	2377	2378	2379	2380	2381	2382	2383	2384	2385	2386	2387	2388	2389	2390	2391	2392	2393	2394	2395	2396	2397	2398	2399	2400	2401	2402	2403	2404	2405	2406	2407	2408	2409	2410	2411	2412	2413	2414	2415	2416	2417	2418	2419	2420	2421	2422	2423	2424	2425	2426	2427	2428	2429	2430	2431	2432	2433	2434	2435	2436	2437	2438	2439	2440	2441	2442	2
--	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	------	---

## ERROR

## 5

A fault-tolerant computer system having an application memory organized as a plurality of cache lines, each cache line being identified by an address in the memory. A FIFO buffer stores a plurality of such cache lines. The system includes at least one CPU for executing instructions stored in the application memory. A checkpoint controller defines a series of repeating checkpoint cycles. The application memory and FIFO buffer are operated under the control of a memory controller. The checkpoint controller also has access to a plurality of registers in the CPU that define the state of that CPU at a point in each checkpoint cycle that is controllable by the controller. When the memory controller receives a cache line from the CPU in response to a write command specifying an address A in the application memory at which the cache line is to be stored, a copy of the cache line as stored in the application memory at A is copied into the FIFO buffer upon receiving the first write command specifying A after the start of the current checkpoint cycle. The cache line received in the write command is then used to overwrite the contents of address A in the application memory. At a predetermined point in each checkpoint cycle, the checkpoint controller causes the CPU to write back to the memory all dirty cache lines; and to store its internal registers defining the state of the CPU in the state memory. The checkpoint controller empties the contents of the FIFO buffer at the end of each checkpoint cycle if no error has been detected by the end of the checkpoint phase of the cycle. If an error is detected, the contents of the FIFO buffer are read back into the application memory, and the contents of the state memory are read back into the CPU. The system is then restarted after completing any hardware configuration needed for the restart.



10




FIGURE 3



## PATENT APPLICATION

DECLARATION AND POWER OF ATTORNEY  
FOR PATENT APPLICATION

ATTORNEY DOCKET NO. 10991625-1

As a below named inventor, I hereby declare that:

My residence/post office address and citizenship are as stated below next to my name;

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled:

CHECKPOINT COMPUTER SYSTEM UTILIZING A FIFO BUFFER TO RE-SYNCHRONIZE AND RECOVER THE SYSTEM ON THE DETECTION OF AN ERROR

the specification of which is attached hereto unless the following box is checked:

☐ was filed on \_\_\_\_\_ as US Application Serial No. or PCT International Application Number \_\_\_\_\_ and was amended on \_\_\_\_\_ (if applicable).

I hereby state that I have reviewed and understood the contents of the above-identified specification, including the claims, as amended by any amendment(s) referred to above. I acknowledge the duty to disclose all information which is material to patentability as defined in 37 CFR 1.56.

## Foreign Application(s) and/or Claim of Foreign Priority

I hereby claim foreign priority benefits under Title 35, United States Code Section 119 of any foreign application(s) for patent or inventor(s) certificate listed below and have also identified below any foreign application for patent or inventor(s) certificate having a filing date before that of the application on which priority is claimed:

COUNTRY	APPLICATION NUMBER	DATE FILED	PRIORITY CLAIMED UNDER 35 U.S.C. 119
			YES: <input type="checkbox"/> NO: <input type="checkbox"/>
			YES: <input type="checkbox"/> NO: <input type="checkbox"/>

## Provisional Application

I hereby claim the benefit under Title 35, United States Code Section 119(e) of any United States provisional application(s) listed below:

APPLICATION SERIAL NUMBER	FILING DATE

## U. S. Priority Claim

I hereby claim the benefit under Title 35, United States Code, Section 120 of any United States application(s) listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States application in the manner provided by the first paragraph of Title 35, United States Code Section 112, 1 acknowledge the duty to disclose material information as defined in Title 37, Code of Federal Regulations, Section 1.56(a) which occurred between the filing date of the prior application and the national or PCT international filing date of this application:

APPLICATION SERIAL NUMBER	FILING DATE	STATUS (patented/pending/abandoned)

## POWER OF ATTORNEY:

As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith:

Customer Number 022879

Place Customer  
Number Bar Code  
Label hereSend Correspondence to:  
HEWLETT-PACKARD COMPANY  
Intellectual Property Administration  
P.O. Box 272400  
Fort Collins, Colorado 80528-9599

Direct Telephone Calls To:

Calvin Ward  
925-855-0413

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

Full Name of Inventor: Zheng Zhang

Citizenship: P.R. China

Residence: 3520 Casabella Court

Post Office Address: San Jose, CA 95148

Inventor's Signature

9/5/00

ATTORNEY DOCKET NO. 10991625-1

**Citizenship:** USA

Post Office Address: Stanford, CA 94305

Manohar K Prabhu  
Inventor's signature

SEPTEMBER 5, 2000  
Date

**Citizenship:**

**Post Office Address:**

**Inventor's signature**

**Date**

**Citizenship:**

**Post Office Address:**

~~Inventor's signature~~

**Date**

**Citizenship:**

**Post Office Address:**

**Inventor's signature**

**Date**

**Citizenship:**

**Post Office Address:**

**Inventor's signature**

**Date**

**Citizenship:**

**Post Office Address:**

Inventor's signature

~ Date

**Citizenship:**

**Post Office Address:**

~~Inventor's signature~~

**Date**